

**REMARKS**

In the non-final Office Action, the Examiner withdrew claims 9-18 from consideration; rejected claims 1-8 and 19-27 under 35 U.S.C. § 101 as allegedly lacking utility; rejected claims 1-8 under 35 U.S.C. § 102(b) as allegedly anticipated by Jindal et al. (U.S. Patent No. 6,324,580); and rejected claims 19-27 under 35 U.S.C. § 103(a) as allegedly unpatentable over Jindal et al. and Narendran et al. (U.S. Patent No. 6,070,191).

By this Amendment, Applicants amend claims 7 and 26 to improve form. Applicants respectfully traverse the Examiner's rejections under 35 U.S.C. §§ 101, 102, and 103. Claims 1-27 remain pending, of which claims 9-18 have been withdrawn from consideration by the Examiner.

*RESTRICTION REQUIREMENT*

In paragraph 2 of the Office Action, the Examiner maintained the restriction requirement even though the non-elected claims have been amended to depend from the elected claims. Applicants respectfully request the rejoinder of the non-elected claims (claims 9-18) upon allowance of the elected claims from which the non-elected claims depend.

*REJECTION UNDER 35 U.S.C. § 101*

In paragraph 4 of the Office Action, the Examiner rejected claims 1-8 and 19-27 under 35 U.S.C. § 101 as allegedly lacking patentable utility. The Examiner alleged that claims 1, 8, 19, 26, and 27 merely identify or define "a data process in which to be manipulated without giving rise to a concrete, useful and tangible result" (Office Action, page 3). Applicants respectfully traverse the rejection and submit that the Examiner's allegation falls short of establishing a prima facie case to deny patentability under 35 U.S.C. § 101.

According to the Interim Guidelines for Examination of Patent Applications for Patent Subject Matter Eligibility, the Examiner must determine whether the claimed invention falls within an enumerated statutory category (Interim Guidelines For Examination of Patent Applications for Patent Subject Matter Eligibility, 1300 OG 142, § IV. B., November 22, 2005).

In this case, each of the claims properly falls within one or more of the enumerated statutory categories. For example, claims 1 and 19 are method claims that can fall within the process category, and claims 7, 8, 26, and 27 are system claims that can fall within the machine category.

The Interim Guidelines indicate that the burden is on the Examiner to set forth a prima facie case of unpatentability; therefore, if the Examiner determines that it is more likely than not that the claimed subject matter falls outside all of the statutory categories, the Examiner must provide an explanation (Id. at §IV. B.). In this case, the Examiner's allegation that the claims identify or define a data process falls short of establishing that the claims fall outside all of the statutory categories, because a "process" is one of the statutory categories.

The Interim Guidelines indicate that if a claimed invention falls outside all of the statutory categories, the claims might still satisfy the requirements of 35 U.S.C. § 101 if the claims, as a whole, are directed to a particular application of an abstract idea, natural phenomenon, or law of nature, rather than to the abstract idea, natural phenomenon, or law of nature itself (Id. at § IV. C.). The Examiner has not even alleged that one or more of Applicants' claims are directed to an abstract idea, natural phenomenon, or a law of nature. Therefore, the Examiner's allegation that the claims identify or define a data process falls short of establishing that the claims are directed to an abstract idea, a natural phenomenon, or a law of nature.

Because the Examiner has not established a prima facie case to deny patentability under 35 U.S.C. § 101, Applicants are not required to show a concrete, useful, and tangible result.

Accordingly, Applicants respectfully request reconsideration and withdrawal of the rejection of claims 1-8 and 19-27 under 35 U.S.C. § 101.

*REJECTION UNDER 35 U.S.C. § 102(b) BASED ON JINDAL ET AL.*

In paragraph 5 of the Office Action, the Examiner rejected claims 1-8 under 35 U.S.C. § 102(b) as allegedly anticipated by Jindal et al. Applicants respectfully traverse the rejection.

A proper rejection under 35 U.S.C. § 102 requires that a single reference teach every aspect of the claimed invention. Any feature not directly taught must be inherently present. In other words, the identical invention must be shown in as complete detail as contained in the claim. See M.P.E.P. § 2131. Jindal et al. does not disclose or suggest the combination of features recited in claims 1-8.

Independent claim 1, for example, is directed to a method for distributing data in a system that includes a plurality of servers. The method comprises identifying ones of the servers to store a replica of the data based on at least one of utilization of the servers, prior data distribution involving the servers, or failure correlation properties associated with the servers; and placing the replicas of the data at the identified servers.

Jindal et al. does not disclose or suggest the combination of features recited in claim 1. For example, Jindal et al. does not disclose or suggest identifying ones of the servers to store a replica of the data based on at least one of utilization of the servers, prior data distribution involving the servers, or failure correlation properties associated with the servers.

The Examiner alleged that Jindal et al. discloses these features and cited the Abstract, the Summary, column 4, lines 40-67, and column 5, lines 57-60, of Jindal et al. for support (Office Action, page 3). Applicants respectfully disagree.

In the Abstract, Jindal et al. discloses a method for load balancing requests for a replicated service or application among a plurality of servers operating instances of the replicated service or application, and choosing a preferred server from the plurality of servers to receive a request. Nowhere in this section, or elsewhere, does Jindal et al. disclose or remotely suggest identifying ones of the servers to store a replica of the data based on at least one of utilization of the servers, prior data distribution involving the servers, or failure correlation properties associated with the servers, as required by claim 1.

In the Summary, Jindal et al. discloses a load balancing policy that specifies one or more factors to be used in determining a server that is to receive a client request (col. 2, lines 35-39). Nowhere in this section, or elsewhere, does Jindal et al. disclose or remotely suggest identifying ones of the servers to store a replica of the data based on at least one of utilization of the servers, prior data distribution involving the servers, or failure correlation properties associated with the servers, as required by claim 1.

At column 4, lines 40-67, Jindal et al. discloses:

In a present embodiment of the invention, information concerning the operation of computer servers executing a replicated service is collected and processed to identify a preferred server (e.g., the server with the smallest load or shortest response time). Illustrative pieces of information that are collected include a server's response time, its distance from a central server (such as a name server providing DNS services), its operational status (e.g., whether it is up or down), etc.

For purposes of the present invention a replicated service is a service (e.g., web browsing, electronic mail) that is available on multiple servers. For example, an organization providing a service or application that is visited or invoked by numerous clients may employ several web servers to handle the requests. Each of the several servers is

considered to operate a separate instance of the replicated service or application. Individual users may thus be routed to, and their requests satisfied by, any of the several servers.

The collected information is then analyzed and a preferred server is identified in accordance with a selected policy. In accordance with one illustrative policy, the preferred server is the server that is least-loaded. Another policy identifies the preferred server as being the closest server. After the preferred server is identified, subsequent requests for the replicated service or application are directed to that server. For example, in a web-browsing environment a DNS lookup table, or zone file, is updated to indicate that requests for the replicated service are to be routed to the preferred server.

In this section, Jindal et al. discloses that information is collected from the servers to identify the preferred server to receive a client request. Nowhere in this section, or elsewhere, does Jindal et al. disclose or remotely suggest identifying ones of the servers to store a replica of the data based on at least one of utilization of the servers, prior data distribution involving the servers, or failure correlation properties associated with the servers, as required by claim 1. Instead, Jindal et al. discloses that a replicated service is simply a service that is available on multiple servers (col. 4, lines 49-51). Jindal et al. does not disclose or suggest the manner in which the replicated service is stored on the servers.

At column 5, lines 57-63, Jindal et al. discloses:

In one embodiment of the present invention, the specific server identified in the zone file is determined according to a selected policy, as discussed below. Further, the server identified in zone file 104 is updated from time to time in accordance with the selected policy in order to distribute client requests among the servers offering the replicated service.

In this section, Jindal et al. discloses that a specific server is determined according to a selected policy. Nowhere in this section, or elsewhere, does Jindal et al. disclose or remotely suggest identifying ones of the servers to store a replica of the data based on at least one of utilization of the servers, prior data distribution involving the servers, or failure correlation properties associated with the servers, as required by claim 1.

Jindal et al. also does not disclose or suggest placing the replicas of the data at the identified servers, as further recited in claim 1. The Examiner alleged that Jindal et al. discloses this feature and cited column 4, lines 40-67, of Jindal et al. for support (Office Action, page 3). Applicants disagree.

Column 4, lines 40-67, of Jindal et al. is reproduced above. In this section, Jindal et al. discloses that information is collected from the servers to identify the preferred server to receive a client request. Nowhere in this section, or elsewhere, does Jindal et al. disclose or remotely suggest placing the replicas of the data at the identified servers, as required by claim 1.

Further, Applicants note that the Examiner admitted with regard to another claim (i.e., claim 19) that Jindal et al. does not disclose replicas (Office Action, page 6). If this is so, Applicants cannot comprehend how the Examiner can allege that Jindal et al. discloses identifying ones of the servers to store a replica of the data or placing the replicas of the data, as required by claim 1.

For at least these reasons, Applicants submit that claim 1 is not anticipated by Jindal et al. Claims 2-6 depend from claim 1 and are, therefore, not anticipated by Jindal et al. for at least the reasons given with regard to claim 1. Claims 2-6 are also not anticipated by Jindal et al. for reasons of their own.

For example, claim 2 recites identifying underutilized ones of the servers as candidates to store the replicas of the data. Jindal et al. does not disclose or suggest the combination of features recited in claim 2.

The Examiner alleged that Jindal et al. discloses the feature of claim 2 and cited the Abstract, the Summary, and column 4, lines 40-67, of Jindal et al. for support (Office Action, page 3). Applicants respectfully disagree.

As noted above, the Abstract of Jindal et al. discloses a method for load balancing requests for a replicated service or application among a plurality of servers operating instances of the replicated service or application, and choosing a preferred server from the plurality of servers to receive a request. Nowhere in this section, or elsewhere, does Jindal et al. disclose or remotely suggest identifying underutilized ones of the servers as candidates to store the replicas of the data, as required by claim 2.

As noted above, the Summary of Jindal et al. discloses a load balancing policy that specifies one or more factors to be used in determining a server that is to receive a client request (col. 2, lines 35-39). Nowhere in this section, or elsewhere, does Jindal et al. disclose or remotely suggest identifying underutilized ones of the servers as candidates to store the replicas of the data, as required by claim 2.

Column 4, lines 40-67, of Jindal et al. is reproduced above. In this section, Jindal et al. discloses that information is collected from the servers to identify the preferred server to receive a client request. Nowhere in this section, or elsewhere, does Jindal et al. disclose or remotely suggest identifying underutilized ones of the servers as candidates to store the replicas of the data, as required by claim 2.

For at least these additional reasons, Applicants submit that claim 2 is not anticipated by Jindal et al. Claim 3 depends from claim 2 and is, therefore, not anticipated by Jindal et al. for at least the additional reasons given with regard to claim 2.

Claim 4 recites identifying ones of the servers that have not been involved in a recent data distribution as candidates to store the replicas of the data. Jindal et al. does not disclose or suggest the combination of features recited in claim 4.

The Examiner alleged that Jindal et al. discloses the feature of claim 4 and cited the Abstract, the Summary, column 4, lines 40-67, and column 5, lines 57-60, of Jindal et al. for support (Office Action, page 4). Applicants respectfully disagree.

As noted above, the Abstract of Jindal et al. discloses a method for load balancing requests for a replicated service or application among a plurality of servers operating instances of the replicated service or application, and choosing a preferred server from the plurality of servers to receive a request. Nowhere in this section, or elsewhere, does Jindal et al. disclose or remotely suggest identifying ones of the servers that have not been involved in a recent data distribution as candidates to store the replicas of the data, as required by claim 4.

As noted above, the Summary of Jindal et al. discloses a load balancing policy that specifies one or more factors to be used in determining a server that is to receive a client request (col. 2, lines 35-39). Nowhere in this section, or elsewhere, does Jindal et al. disclose or remotely suggest identifying ones of the servers that have not been involved in a recent data distribution as candidates to store the replicas of the data, as required by claim 4.

Column 4, lines 40-67, of Jindal et al. is reproduced above. In this section, Jindal et al. discloses that information is collected from the servers to identify the preferred server to receive a client request. Nowhere in this section, or elsewhere, does Jindal et al. disclose or remotely suggest identifying ones of the servers that have not been involved in a recent data distribution as candidates to store the replicas of the data, as required by claim 4.



Column 5, lines 57-63, of Jindal et al. is reproduced above. In this section, Jindal et al. discloses that a specific server is determined according to a selected policy. Nowhere in this section, or elsewhere, does Jindal et al. disclose or remotely suggest identifying ones of the servers that have not been involved in a recent data distribution as candidates to store the replicas of the data, as required by claim 4.

For at least these additional reasons, Applicants submit that claim 4 is not anticipated by Jindal et al.

Claim 5 recites identifying system conditions that affect two or more of the servers, and identifying ones of the servers as candidates to store the replicas of the data based on the identified system conditions. Jindal et al. does not disclose or suggest the combination of features recited in claim 5.

The Examiner alleged that Jindal et al. discloses the features of claim 5 and cited column 4, lines 40-67, column 5, lines 57-60, and column 6, lines 32-46, of Jindal et al. for support (Office Action, page 4). Applicants respectfully disagree.

Column 4, lines 40-67, of Jindal et al. is reproduced above. In this section, Jindal et al. discloses that information is collected from the servers to identify the preferred server to receive a client request. Nowhere in this section, or elsewhere, does Jindal et al. disclose or remotely suggest identifying system conditions that affect two or more of the servers, or identifying ones of the servers as candidates to store the replicas of the data based on the identified system conditions, as required by claim 5.

Column 5, lines 57-63, of Jindal et al. is reproduced above. In this section, Jindal et al. discloses that a specific server is determined according to a selected policy. Nowhere in this

section, or elsewhere, does Jindal et al. disclose or remotely suggest identifying system conditions that affect two or more of the servers, or identifying ones of the servers as candidates to store the replicas of the data based on the identified system conditions, as required by claim 5.

At column 6, lines 32-46, Jindal et al. discloses:

In this mode of operation, status objects 200a, 200b and 200c are invoked on nameserver 100 for the purpose of gathering information from servers 110, 112 and 114, respectively. The configuration and purpose of the status objects depend upon the policy that has been selected for choosing a preferred server. For example, where the selected policy requires choosing the least-loaded server (e.g., that which has the fastest response time), each status object measures the response time of its associated server. Illustratively, this may be accomplished by issuing a Ping (or similar) command to the server and measuring the response time. As another example, where the selected policy requires choosing the closest server the status object is illustratively configured to measure the number of hops from nameserver 100 to the object's associated server.

In this section, Jindal et al. discloses choosing a least-loaded server or a closest server. Jindal et al. explicitly discloses that this server is selected to receive a client request (col. 2, lines 36-43). Nowhere in this section, or elsewhere, does Jindal et al. disclose or remotely suggest identifying system conditions that affect two or more of the servers, or identifying ones of the servers as candidates to store the replicas of the data based on the identified system conditions, as required by claim 5.

For at least these additional reasons, Applicants submit that claim 5 is not anticipated by Jindal et al.

Claim 6 recites that a number of the replicas of the data stored by the servers is user-configurable. Jindal et al. does not disclose or suggest the combination of features recited in claim 6.

The Examiner alleged that Jindal et al. discloses the feature of claim 6 and cited the Abstract, the Summary, column 4, lines 40-67, and column 5, lines 57-60, of Jindal et al. for support (Office Action, page 4). Applicants respectfully disagree.

As noted above, the Abstract of Jindal et al. discloses a method for load balancing requests for a replicated service or application among a plurality of servers operating instances of the replicated service or application, and choosing a preferred server from the plurality of servers to receive a request. Nowhere in this section, or elsewhere, does Jindal et al. disclose or remotely suggest that a number of the replicas of the data stored by the servers is user-configurable, as required by claim 6.

As noted above, the Summary of Jindal et al. discloses a load balancing policy that specifies one or more factors to be used in determining a server that is to receive a client request (col. 2, lines 35-39). Nowhere in this section, or elsewhere, does Jindal et al. disclose or remotely suggest that a number of the replicas of the data stored by the servers is user-configurable, as required by claim 6.

Column 4, lines 40-67, of Jindal et al. is reproduced above. In this section, Jindal et al. discloses that information is collected from the servers to identify the preferred server to receive a client request. Nowhere in this section, or elsewhere, does Jindal et al. disclose or remotely suggest that a number of the replicas of the data stored by the servers is user-configurable, as required by claim 6.

Column 5, lines 57-63, of Jindal et al. is reproduced above. In this section, Jindal et al. discloses that a specific server is determined according to a selected policy. Nowhere in this

section, or elsewhere, does Jindal et al. disclose or remotely suggest that a number of the replicas of the data stored by the servers is user-configurable, as required by claim 6.

For at least these additional reasons, Applicants submit that claim 6 is not anticipated by Jindal et al.

Independent claims 7 and 8 recite features similar to (yet possibly different in scope from) features recited in claim 1. Claims 7 and 8 are, therefore, not anticipated by Jindal et al. for at least reasons similar to reasons given with regard to claim 1.

Accordingly, Applicants respectfully request reconsideration and withdrawal of the rejection of claims 1-8 under 35 U.S.C. § 102(b) based on Jindal et al.

*REJECTION UNDER 35 U.S.C. § 103(a) BASED ON  
JINDAL ET AL. AND NARENDRAN ET AL.*

In paragraph 6 of the Office Action, the Examiner rejected claims 19-27 under 35 U.S.C. § 103(a) as allegedly unpatentable over Jindal et al. in view of Narendran et al. Applicants respectfully traverse the rejection.

Independent claim 19 is directed to a method for redistributing chunks of data in a system that includes a plurality of servers that store replicas of the chunks. The method comprises monitoring utilization of the servers; determining whether to redistribute any of the replicas; selecting one or more of the replicas to redistribute based on the utilization of the servers; selecting one or more of the servers to which to move the one or more replicas; and moving the one or more replicas to the selected one or more servers.

Neither Jindal et al. nor Narendran et al., whether taken alone or in any reasonable combination, discloses or suggests the combination of features recited in claim 19. For example,

Jindal et al. and Narendran et al. do not disclose or suggest determining whether to redistribute any of the replicas.

The Examiner alleged that Jindal et al. discloses this feature and cited column 6, lines 31-45 and 56-64, of Jindal et al. for support (Office Action, page 5). Applicants respectfully disagree.

Column 6, lines 31-45, of Jindal et al. is reproduced above. In this section, Jindal et al. discloses choosing a least-loaded server or a closest server. Jindal et al. explicitly discloses that this server is selected to receive a client request (col. 2, lines 36-43). Nowhere in this section, or elsewhere, does Jindal et al. disclose or remotely suggest determining whether to redistribute any of the replicas, as required by claim 19.

The Examiner admitted that Jindal et al. does not disclose replicas (Office Action, page 6). If this is so, Applicants cannot comprehend how the Examiner can allege that Jindal et al. discloses determining whether to redistribute any of the replicas, as required by claim 19.

Nevertheless, the Examiner alleged that Narendran et al. discloses data distribution techniques for load-balanced fault-tolerant web access including redistributing replicated data from a failed server to achieve rebalance and cited column 12, lines 12-38, of Narendran et al. for support (Office Action, page 6). Applicants submit that the disclosure of Narendran et al. does not support the Examiner's allegation and, in fact, teaches away from the features recited in claim 19.

At column 12, lines 12-38, Narendran et al. discloses:

The load balance that is achieved through initial distribution of the documents may be disturbed during the operation of the server cluster due to the following changes: (1) a server failure; (2) changes in the access probabilities of the documents, i.e., some documents which were initially "hot" could become "cold" and vice-versa, and (3) changes in the capacity of the document servers. In the case of dedicated web servers, the

first and second changes are generally more likely than the third. In the event any of these changes occur, it is desirable to be able to rebalance the server loads without any major reconfiguration of the system. More specifically, the rebalancing should be accomplished without redistributing documents among the servers, since this would generally involve overhead in moving the documents between servers and may affect the availability of the system. It is preferable to instead achieve rebalance by adjusting only the redirection probabilities used by the redirection server. As will be shown below, this rebalancing can be characterized as a network flow problem. It should be noted that the network flow approach is suitable for use in situations in which documents on a failed server are replicated and therefore available on another server or servers. In situations in which the documents on the failed server are not replicated, the above-described binning algorithm may be used to redistribute the documents from the failed server to achieve rebalance with a minimal amount of document movement.

In this section, Narendran et al. discloses rebalancing server loads by adjusting the redirection probabilities used by the redirection server when a server failure, changes in the access probabilities of the documents, or changes in the capacity of the document servers occurs. Nowhere in this section, or elsewhere, does Narendran et al. disclose or remotely suggest determining whether to redistribute any of the replicas, as required by claim 19. Instead, Narendran et al. discloses "the rebalancing should be accomplished without redistributing documents among the servers, since this would generally involve overhead in moving the documents between servers and may affect the availability of the system" (col. 12, lines 23-26). Accordingly, Narendran et al. specifically teaches away from redistributing documents (which the Examiner apparently equates to the replicas recited in claim 19).

Narendran et al. also discloses that in situations in which the documents on the failed server are not replicated, the algorithm could be used to redistribute the documents from the failed server to achieve rebalance with a minimal amount of document movement (col. 12, lines 33-38). In this situation, however, the documents cannot reasonably be interpreted as replicas since Narendran et al. specifically discloses that the documents are not replicated (col. 12, lines 33-35).

At column 6, lines 56-64, Jindal et al. discloses:

Illustratively, for each replicated service (or application) that is to be monitored (i.e., that is subject to load balancing) on a server, a separate status object operates on nameserver 100. In addition, each status object illustratively performs a single function (e.g., determine response time, determine a server's distance from nameserver 100). In alternative embodiments of the invention, however, a single status object may monitor multiple servers or services and/or perform multiple functions.

In this section, Jindal et al. discloses a status object for each replicated service or application to be monitored. Nowhere in this section, or elsewhere, does Jindal et al. disclose or remotely suggest determining whether to redistribute any of the replicas, as required by claim 19.

Jindal et al. and Narendran et al. also do not disclose or suggest selecting one or more replicas to redistribute based on the utilization of the servers, as further recited in claim 19. The Examiner alleged that Jindal et al. discloses this feature and cited column 6, lines 31-45 and 56-64, of Jindal et al. for support (Office Action, page 6). Applicants disagree.

Column 6, lines 31-45 and 56-64, of Jindal et al. are reproduced above. Nowhere in these sections, or elsewhere, for reasons similar to reasons given above, does Jindal et al. disclose or remotely suggest selecting one or more replicas to redistribute based on the utilization of the servers, as required by claim 19. Narendran et al. also does not disclose or suggest this feature for reasons similar to reasons given above.

Because Jindal et al. and Narendran et al., whether taken alone or in any reasonable combination, do not disclose or suggest determining whether to redistribute any of the replicas or selecting one or more of the replicas to redistribute based on the utilization of the servers, Jindal et al. and Narendran et al. cannot disclose or suggest selecting one or more of the servers to which to move the one or more replicas or moving the one or more replicas to the selected one or more servers, as further recited in claim 19.

For at least these reasons, Applicants submit that claim 19 is patentable over Jindal et al. and Narendran et al., whether taken alone or in any reasonable combination. Claims 20-25 depend from claim 19 and are, therefore, patentable over Jindal et al. and Narendran et al. for at least the reasons given with regard to claim 19.

Independent claims 26 and 27 recite features similar to (but possibly different in scope from) features recited in claim 19. Claims 26 and 27 are, therefore, patentable over Jindal et al. and Narendran et al., whether taken alone or in any reasonable combination, for at least reasons similar to reasons given with regard to claim 19.

Accordingly, Applicants respectfully request reconsideration and withdrawal of the rejection of claims 19-27 under 35 U.S.C. § 103(a) based on Jindal et al. and Narendran et al.

#### *CONCLUSION*

In view of the foregoing amendments and remarks, Applicants respectfully request the Examiner's reconsideration of the application and the timely allowance of pending claims 1-27.

As Applicants' remarks with respect to the Examiner's rejections overcome the rejections, Applicants' silence as to certain assertions by the Examiner in the Office Action or certain requirements that may be applicable to such rejections (e.g., whether a reference constitutes prior art, motivation to combine references, etc.) is not a concession by Applicants that such assertions are accurate or that such requirements have been met, and Applicants reserve the right to dispute these assertions/requirements in the future.

If the Examiner believes that the application is not now in condition for allowance, Applicants respectfully request that the Examiner contact the undersigned to discuss any outstanding issues.



To the extent necessary, a petition for an extension of time under 37 C.F.R. § 1.136 is hereby made. Please charge any shortage in fees due in connection with the filing of this paper, including extension of time fees, to Deposit Account No. 50-1070 and please credit any excess fees to such deposit account.

Respectfully submitted,

HARRITY SNYDER, L.L.P.

/Paul A. Harrity/  
Paul A. Harrity  
Reg. No. 39,574

Date: August 18, 2006

11350 Random Hills Road  
Suite 600  
Fairfax, Virginia 22030  
(571) 432-0800